arXiv:1512.06527v2 [math.NA] 10 Nov 2016

Towards tensor-based methods for the numerical approximation of the Perron–Frobenius and Koopman operator

Stefan Klus[1] and Christof Schütte[1,2]

[1]Department of Mathematics and Computer Science, Freie Universität Berlin, Germany
[2]Zuse Institute Berlin, Germany

Abstract

The global behavior of dynamical systems can be studied by analyzing the eigenvalues and corresponding eigenfunctions of linear operators associated with the system. Two important operators which are frequently used to gain insight into the system's behavior are the Perron–Frobenius operator and the Koopman operator. Due to the curse of dimensionality, computing the eigenfunctions of high-dimensional systems is in general infeasible. We will propose a tensor-based reformulation of two numerical methods for computing finite-dimensional approximations of the aforementioned infinite-dimensional operators, namely Ulam's method and Extended Dynamic Mode Decomposition (EDMD). The aim of the tensor formulation is to approximate the eigenfunctions by low-rank tensors, potentially resulting in a significant reduction of the time and memory required to solve the resulting eigenvalue problems, provided that such a low-rank tensor decomposition exists. Typically, not all variables of a high-dimensional dynamical system contribute equally to the system's behavior, often the dynamics can be decomposed into slow and fast processes, which is also reflected in the eigenfunctions. Thus, the weak coupling between different variables might be approximated by low-rank tensor cores. We will illustrate the efficiency of the tensor-based formulation of Ulam's method and EDMD using simple stochastic differential equations.

1 Introduction

The Perron–Frobenius operator and the Koopman operator enable the analysis of the global behavior of dynamical systems. Eigenfunctions of these operators can be used to extract the dominant dynamics, to detect almost invariant sets, or to decompose the system into fast and slow processes [1, 2, 3, 4, 5, 6, 7]. Assume the state space of your system is $\mathbb{R}^d$ and you want to discretize each direction using n grid points (or boxes if Ulam's method is used), then overall n^d values need to be stored. Even for $d = 10$ and $n = 10$, more than 70 gigabyte of storage space would be required, whereas typical systems might have hundreds or thousands of dimensions and naturally also require a more fine-grained discretization. This so-called *curse of dimensionality* can be overcome by using tensor formats which compress the data and store only the information that is relevant for the reconstruction. In general, only an approximation of the original data can be retrieved. Approximating the objects under consideration by sums of low-rank tensor products has become a powerful approach for tackling high-dimensional problems [8] and in many physically significant problems near-linear complexity can be achieved since the separation rank depends only weakly on the dimension [9]. For high-dimensional systems exhibiting multiscale behavior, it might be possible to represent the weak coupling between different variables by low-rank tensor cores. The leading eigenfunctions of the Koopman operator, for instance, are typically almost constant for the fast variables of the system and depend mainly on the slowly changing variables (see [6]). Thus, using tensor-based algorithms could reduce the amount of time and memory required to compute and store eigenfunctions significantly. In this way, analyzing high-dimensional systems that could not be tackled using standard methods might become feasible.

Tensors, in our sense, are just multidimensional arrays as shown in Figure 1. Here and in what follows, standard vectors will be denoted by lower-case letters, e.g. v, matrices by upper-case letters, e.g. A, and tensors by the corresponding bold symbols, e.g. $\mathbf{x}$. It is important to note that tensors are typically not explicitly given – for example by observed data –, but only implicitly as solutions of systems of linear or nonlinear equations or eigenvalue

problems [10]. Thus, numerical methods that operate directly on tensor approximations need to be developed since the full tensors cannot be stored or handled anymore in practice.

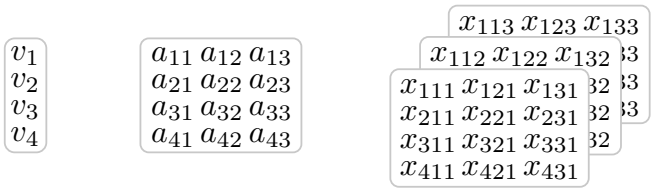

Figure 1: Tensors as multidimensional arrays. Here, $v = (v_i) \in \mathbb{R}^4$, $A = (a_{ij}) \in \mathbb{R}^{4\times 3}$, and $\mathbf{x} = (x_{ijk}) \in \mathbb{R}^{4\times 3\times 3}$.

Over the last years, low-rank tensor approximation approaches have become increasingly popular in the scientific computing community and are now becoming a standard tool to cope with large-scale problems that could not be handled before by standard numerical methods. An overview of different low-rank tensor approximation approaches can be found in [10].

In this paper, we will show that the use of low-rank tensor approximation schemes potentially enables the computation of eigenfunctions of high-dimensional systems. The aim of this paper, however, is not to show that our approach is more efficient – the tensor algorithms are mainly implemented in Matlab, a comparison with highly optimized numerical libraries implemented in C or C++ would not lead to meaningful results, developing high performance libraries for large-scale tensor problems is a separate task –, but to derive a tensor-based reformulation of existing methods and to show equivalency so that the theory available for the conventional matrix-vector based formulation can be carried over to multi-dimensional arrays. Furthermore, tensors could enable low-rank approximations of the Perron–Frobenius and Koopman operator as well as their eigenfunctions. One of the main future goals is to combine low-rank tensor decomposition techniques and the splitting of the dynamics into fast and slow processes. In [6], it has been shown that such a splitting of a multi-scale system exists and can be exploited to extract the slow dynamics. Another open problem is the generation of the low-rank approximations of the operators. Currently, the canonical tensor format representations are converted to the tensor-train format, which is time-consuming and in general leads to high ranks. Ideally, low-rank TT approximations should be directly generated from the given data.

We will start by introducing standard methods such as Ulam's method and the recently developed Extended Dynamic Mode Decomposition (EDMD) to approximate the eigenfunctions of the Perron–Frobenius operator and the Koopman operator in Section 2. Then, in Section 3, different tensor formats will be described. In Section 4, we will reformulate Ulam's method and EDMD as tensor-based methods. Section 5 contains a brief summary of simple power iteration schemes for the resulting tensor-based (generalized) eigenvalue problems. Simple examples which illustrate the proposed approaches are shown in Section 6. In Section 7, we will conclude with a short summary and possible future work.

2 Perron–Frobenius and Koopman operator approximation

In this section, we will briefly introduce the Perron–Frobenius operator $\mathcal{P}$ and the Koopman operator $\mathcal{K}$ as well as numerical methods to compute finite-dimensional approximations, namely Ulam's method and EDMD. The main difference between Ulam's method and EDMD is that the former uses indicator functions[1] for a given box discretization of the domain while the latter allows arbitrary ansatz functions such as monomials, Hermite polynomials, trigonometric functions, or radial basis functions. Although EDMD was primarily developed for the approximation of the Koopman operator, it can be used to compute eigenfunctions of the Perron–Frobenius operator as well [12]. Analogously, Ulam's method can also be used to compute eigenfunctions of the Koopman operator.

[1] Higher-order methods for the approximation of the Perron–Frobenius operator have been proposed in [11].

2.1 Perron–Frobenius and Koopman operator

Let $S : \mathcal{X} \to \mathcal{X}$ be a dynamical system defined on a domain $\mathcal{X}$, for example $\mathcal{X} \subseteq \mathbb{R}^d$. Then the Perron–Frobenius operator or transfer operator $\mathcal{P}$ is defined by

$$\int g \cdot \mathcal{P} f \, dm = \int (g \circ S) \cdot f \, dm, \tag{1}$$

for all $f, g \in \mathcal{F}$, where $\mathcal{F}$ is an appropriately defined function space and $\circ$ denotes function composition. We assume in what follows that $\mathcal{F} = L^2(\mathcal{X})$. The aim is to compute eigenfunctions of the Perron–Frobenius operator, given by

$$\mathcal{P}\varphi_i = \lambda_i \varphi_i.$$

The eigenfunction φ_1 corresponding to $\lambda_1 = 1$ is the invariant density of the system, i.e. $\mathcal{P}\varphi_1 = \varphi_1$. The magnitude of the second largest eigenvalue λ_2 can be interpreted as the rate at which initial densities converge to the invariant density (for more details and assumptions about the dynamical system, see e.g. [6] and references therein). More generally, the leading eigenvalues of the Perron–Frobenius operator close to one correspond to the slowly converging transients of the system. The Koopman operator $\mathcal{K}$, on the other hand, is defined by

$$\mathcal{K} f = f \circ S$$

and acts on functions $f : \mathcal{X} \to \mathbb{C}$, $f \in \mathcal{F}$. Correspondingly, the stochastic Koopman operator is defined by $\mathcal{K} f = \mathbb{E}[\, f \circ S \,]$, where $\mathbb{E}[\,\cdot\,]$ denotes the expected value with respect to the probability measure underlying $S(x)$. We will only introduce the required notation and focus mainly on discrete-time dynamical systems, for more details on the Koopman operator and its properties, we refer to [4, 5, 13]. While the Perron–Frobenius operator describes the evolution of *densities*, the Koopman operator describes the evolution of *observables*, which could be measurements or sensor probes [4]. Instead of analyzing orbits $\{x,\, S(x),\, S^2(x),\, \dots\}$ of the dynamical system, we now analyze the measurements $\{f(x),\, f(S(x)),\, f(S^2(x)),\, \dots\}$ at these points.

The Koopman operator $\mathcal{K}$ is the adjoint of the Perron–Frobenius operator $\mathcal{P}$ and thus an infinite-dimensional but linear operator. A finite-dimensional approximation (computed using generalized Galerkin methods) of this operator captures the dynamics of a nonlinear dynamical system without necessitating a linearization around a fixed point [4, 5]. We are again interested in eigenfunctions of the operator, given by

$$\mathcal{K}\varphi_i = \lambda_i \varphi_i.$$

Let $f : \mathcal{X} \to \mathbb{R}$ be an observable of the system that can be written as a linear combination of the linearly independent eigenfunctions φ_i, i.e.

$$f(x) = \sum_i c_i \varphi_i(x),$$

with $c_i \in \mathbb{C}$. Then

$$(\mathcal{K} f)(x) = \sum_i \lambda_i c_i \varphi_i(x).$$

Analogously, for vector-valued functions $F = [f_1,\, \dots,\, f_n]^T$, we obtain

$$\mathcal{K} F = \begin{bmatrix} \sum_i \lambda_i c_{i,1} \varphi_i \\ \vdots \\ \sum_i \lambda_i c_{i,n} \varphi_i \end{bmatrix} = \sum_i \lambda_i \varphi_i \begin{bmatrix} c_{i,1} \\ \vdots \\ c_{i,n} \end{bmatrix} = \sum_i \lambda_i \varphi_i v_i,$$

where $v_i = [c_{i,1},\, \dots,\, c_{i,n}]^T$. These vectors v_i corresponding to the eigenfunctions φ_i are called Koopman modes.

2.2 Ulam's method

A frequently used method to compute an approximation of the Perron–Frobenius operator is Ulam's method, see e.g. [2, 14, 15, 6]. First, the state space $\mathcal{X}$ is covered by a finite number of disjoint boxes $\{\mathcal{B}_1,\, \dots,\, \mathcal{B}_k\}$. Let $\mathbb{1}_{\mathcal{B}_i}$ be the indicator function for box $\mathcal{B}_i$, i.e.

$$\mathbb{1}_{\mathcal{B}_i}(x) = \begin{cases} 1, & \text{if } x \in \mathcal{B}_i, \\ 0, & \text{otherwise.} \end{cases}$$

Then a finite-dimensional approximation of the operator can be obtained as follows: Using definition (1) leads to

$$\int \mathbb{1}_{\mathcal{B}_j} \cdot \mathcal{P}\mathbb{1}_{\mathcal{B}_i}\, dm = \int (\mathbb{1}_{\mathcal{B}_j} \circ S) \cdot \mathbb{1}_{\mathcal{B}_i}\, dm = \int \mathbb{1}_{S^{-1}(\mathcal{B}_j)} \cdot \mathbb{1}_{\mathcal{B}_i}\, dm = m(S^{-1}(\mathcal{B}_j) \cap \mathcal{B}_i).$$

This relationship can be represented by a matrix $\hat{P} = (\hat{p}_{ij}) \in \mathbb{R}^{k\times k}$ with

$$\hat{p}_{ij} = \frac{m\left(S^{-1}(\mathcal{B}_j) \cap \mathcal{B}_i\right)}{m(\mathcal{B}_i)}.$$

Here, in order to avoid confusion with the approximation $\mathbf{P}$ on a tensor space introduced below, we denote the matrix representation $\hat{P}$ instead of P. The denominator $m(\mathcal{B}_i)$ normalizes the entries $\hat{p}_{ij}$ so that $\hat{P}$ becomes a row-stochastic matrix and defines a finite Markov chain. The left eigenvector corresponding to the eigenvalue $\lambda_1 = 1$ approximates the invariant measure of the Perron–Frobenius operator $\mathcal{P}$.

The entries $\hat{p}_{ij}$ of the matrix $\hat{P}$ represent the probabilities of points being mapped from box $\mathcal{B}_i$ to box $\mathcal{B}_j$ by the dynamical system S. These entries can be estimated by randomly choosing a large number of test points $x_i^{(l)}$, $l = 1, \ldots, n$, in each box $\mathcal{B}_i$ and by counting how many times test points were mapped from box $\mathcal{B}_i$ to box $\mathcal{B}_j$ by S, i.e.

$$\hat{p}_{ij} \approx \frac{1}{n} \sum_{l=1}^{n} \mathbb{1}_{\mathcal{B}_j}\left(S\left(x_i^{(l)}\right)\right). \tag{2}$$

The eigenfunctions of the Perron–Frobenius operator are then approximated by the left eigenvectors of the matrix $\hat{P}$, the eigenfunctions of the Koopman operator by the right eigenvectors.

2.3 Extended dynamic mode decomposition

An approximation of the Koopman operator, the Koopman eigenvalues, eigenfunctions, and eigenmodes can be computed using EDMD. The method requires data, i.e. a set of values x_i and the corresponding $y_i = S(x_i)$ values, $i = 1, \ldots, m$, written in matrix form as

$$X = \begin{bmatrix} x_1 & x_2 & \cdots & x_m \end{bmatrix} \quad \text{and} \quad Y = \begin{bmatrix} y_1 & y_2 & \cdots & y_m \end{bmatrix},$$

and additionally a set of ansatz functions or observables

$$\mathcal{D} = \{\psi_1, \psi_2, \ldots, \psi_k\},$$

with $\psi_i : \mathcal{X} \to \mathbb{R}$. Thus, $X, Y \in \mathbb{R}^{d\times m}$. The vectors x_i are used as collocation points to approximate the integrals required for the approximation of the Koopman operator. Let

$$\Psi = \begin{bmatrix} \psi_1 & \psi_2 & \cdots & \psi_k \end{bmatrix}^T, \tag{3}$$

$\Psi : \mathcal{X} \to \mathbb{R}^k$, be the vector of all ansatz functions, then $\mathcal{K}$ can be approximated by a matrix $\hat{K} \in \mathbb{R}^{k\times k}$, with

$$\hat{K}^T = \hat{A}\hat{G}^+,$$

where $^+$ denotes the pseudoinverse. The matrices $\hat{A}$, $\hat{G} \in \mathbb{R}^{k\times k}$ are defined as

$$\begin{aligned} \hat{A} &= \frac{1}{m} \sum_{l=1}^{m} \Psi(y_l)\Psi(x_l)^T, \\ \hat{G} &= \frac{1}{m} \sum_{l=1}^{m} \Psi(x_l)\Psi(x_l)^T. \end{aligned} \tag{4}$$

As before, we use the $\hat{}$ symbol to distinguish the matrices from the tensor approximations that will be introduced in Section 4. An approximation of the eigenfunction φ_i of the Koopman operator $\mathcal{K}$ is then given by

$$\varphi_i = \xi_i \Psi,$$

where ξ_i is the i-th left eigenvector of the matrix $\hat{K}^T$. Alternatively, the generalized eigenvalue problem

$$\xi_i \hat{A} = \lambda_i \xi_i \hat{G} \tag{5}$$

can be solved, provided that $\hat{G}$ is regular. To compute eigenfunctions of the Perron–Frobenius operator, the corresponding eigenvalue problem

$$\xi_i \hat{A}^T = \lambda_i \xi_i \hat{G}$$

needs to be solved. Note that this formulation is similar to the variational approach to compute eigenfunctions of transfer operators of reversible processes presented in [16, 17]. For more details, we refer the reader to [12].

3 Tensor formats

Several different tensor formats have been developed in the past, e.g. the canonical format, the Tucker format, and the tensor-train format. In this section, we will briefly introduce tensors and the required notation. The overall goal is to rewrite the methods presented in the previous section as tensor-based methods and to take advantage of low-rank tensor approximations and the fact that the dynamics of high-dimensional systems can often be decomposed.

3.1 Full format

A tensor in full format is simply a multidimensional array $\mathbf{v} \in \mathbb{R}^{k_1 \times \cdots \times k_d}$. (A variation of this format is the sparse format which stores only the nonzero entries and is used, for example, in the sparse grid approach [18].) The entries of a tensor $\mathbf{v}$ are indexed by $\mathbf{v}_\mathbf{i} = \mathbf{v}_{i_1,\ldots,i_d} = \mathbf{v}[\mathbf{i}] = \mathbf{v}[i_1, \ldots, i_d]$, where $\mathbf{i} = (i_1, \ldots, i_d)$ is a multi-index. Addition and subtraction are trivially defined element-wise. Multiplication of a tensor $\mathbf{v}$ by a scalar $c \in \mathbb{R}$ is naturally generalized as $(c\mathbf{v})[i_1, \ldots, i_d] = c\mathbf{v}[i_1, \ldots, i_d]$. Matrix-vector multiplication is defined as follows: Given a linear operator $\mathbf{A}$ defined on a tensor space $\mathbb{R}^{k_1 \times \cdots \times k_d}$, with

$$\mathbf{A} = \mathbf{A}[i_1, \ldots, i_d, j_1, \ldots, j_d] \in \mathbb{R}^{k_1 \times \cdots \times k_d \times k_1 \times \cdots \times k_d},$$

the product of $\mathbf{A}$ and $\mathbf{v}$ is

$$(\mathbf{A}\mathbf{v})[i_1, \ldots, i_d] = \sum_{j_1=1}^{k_1} \cdots \sum_{j_d=1}^{k_d} \mathbf{A}[i_1, \ldots, i_d, j_1, \ldots, j_d]\mathbf{v}[j_1, \ldots, j_d]$$

or in shorthand notation, using multi-indices,

$$(\mathbf{A}\mathbf{v})_\mathbf{i} = \sum_\mathbf{j} \mathbf{A}_\mathbf{ij}\mathbf{v}_\mathbf{j}.$$

Furthermore, the inner product of two tensors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^{k_1 \times \cdots \times k_d}$ is defined as

$$\langle \mathbf{v}, \mathbf{w} \rangle = \sum_{i_1=1}^{k_1} \cdots \sum_{i_d=1}^{k_d} \mathbf{v}[i_1, \ldots, i_d]\mathbf{w}[i_1, \ldots, i_d]$$

and the outer product $\mathbf{v} \otimes \mathbf{w} \in \mathbb{R}^{k_1 \times \cdots \times k_d \times k_1 \times \cdots \times k_d}$ as a tensor with entries

$$(\mathbf{v} \otimes \mathbf{w})[i_1, \ldots, i_d, j_1, \ldots, j_d] = \mathbf{v}[i_1, \ldots, i_d]\mathbf{w}[j_1, \ldots, j_d].$$

The outer product $\mathbf{v} \otimes \mathbf{w}$ can be regarded as a linear map that acts on tensors $\mathbb{R}^{k_1 \times \cdots \times k_d}$. Often it is required or convenient to rewrite a tensor as a vector. The *vectorization* of a tensor, denoted $\mathrm{vec}(\mathbf{v})$, where $\mathrm{vec} : \mathbb{R}^{k_1 \times \cdots \times k_d} \to \mathbb{R}^{k_1 \cdots k_d}$, reorders the entries of $\mathbf{v}$ into one column vector. For $\mathbf{v} \in \mathbb{R}^{2 \times 3 \times 2}$, for example,

$$\mathrm{vec}(\mathbf{v}) = \begin{bmatrix} v_{111} & v_{211} & v_{121} & v_{221} & \ldots & v_{122} & v_{222} & v_{132} & v_{232} \end{bmatrix}^T.$$

For our purposes, we will mainly be interested in eigenvalue problems of the form

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{v} \quad \text{or} \quad \mathbf{A}\mathbf{v} = \lambda\mathbf{B}\mathbf{v}.$$

The treatment of tensors in the full format often leads to storage problems. Thus, different formats have been developed to overcome this problem. Instead of working with the full format, we will use compressed formats such as the r-term or TT format for numerical computations in order to minimize computational costs as well as storage requirements. Except for very particular examples, it is impossible to compress the data without any compression error [18]. Typically, the tensor representation is just an approximation of the original data. Below, we will describe different compressed tensor formats, the introduction is based on [18].

3.2 Canonical format

A tensor space is given by $\mathbf{V} = \bigotimes_{\mu=1}^{d} V_\mu$, where $V_1, \dots, V_d$ are vector spaces defined over the same field $\mathbb{K}$, typically $\mathbb{R}$ or $\mathbb{C}$. An *elementary tensor* is defined to be a product of the form

$$\mathbf{v} = v_1 \otimes \cdots \otimes v_d,$$

with $v_\mu \in V_\mu$, $\mu = 1, \dots, d$. An *algebraic tensor* is then a linear combination of elementary tensors, i.e.

$$\mathbf{v} = \sum_{l=1}^{r} v_1^{(l)} \otimes \cdots \otimes v_d^{(l)}.$$

This format is also called r-term format or CP format. That is, instead of trying to store the tensor in the dense format, one only considers tensors that can be written as products of the form $\mathbf{v}[i_1, \dots, i_d] = v_1[i_1] \cdots v_d[i_d]$. If the best approximation of this form is not good enough, then the natural extension is to consider $\mathbf{v}[i_1, \dots, i_d] = \sum_{l=1}^{r} v_1^{(l)}[i_1] \cdots v_d^{(l)}[i_d]$, cf. [19, 9]. Defining

$$\mathcal{R}_r = \left\{ \sum_{l=1}^{r} v_1^{(l)} \otimes \cdots \otimes v_d^{(l)} : v_\mu^{(l)} \in V_\mu \right\}$$

for $r \in \mathbb{N}_0$, which implies that $\{0\} = \mathcal{R}_0 \subset \mathcal{R}_1 \subset \cdots \subset \mathbf{V}$, we call

$$\operatorname{rank}(\mathbf{v}) = \min\{r : \mathbf{v} \in \mathcal{R}_r\}$$

the *tensor rank* of $\mathbf{v}$.

Example 3.1. Let us consider a system of the form $dx_t = -\nabla_x V(x_t)\, dt + \sigma\, dW_t$, where V is the energy landscape associated with the system. The invariant density of this process is given by $\mu(x) = \frac{1}{Z} e^{-\beta V(x)}$, where β and Z are constants [20]. Assume that the state space is three-dimensional and that the potential function can be written as $V(x) = V_1(x_1) + V_2(x_2) + V_3(x_3)$, then

$$\mu(x) = \frac{1}{Z} e^{-\beta V_1(x_1)} e^{-\beta V_2(x_2)} e^{-\beta V_3(x_3)}.$$

Thus, storing the invariant density for a grid with k grid points in each direction in the full format would require an array of size k^3 while storing it in the canonical tensor format would require only a tensor of rank 1 and thus an array of size $3k$. △

Given tensors in the r-term format, basic operations are defined as follows [18]:

- **Addition:**

$$\mathbf{v} = \sum_{l=1}^{r_v} \bigotimes_{\mu=1}^{d} v_\mu^{(l)}, \quad \mathbf{w} = \sum_{l=1}^{r_w} \bigotimes_{\mu=1}^{d} w_\mu^{(l)} \quad \Rightarrow \quad \mathbf{x} = \mathbf{v} + \mathbf{w} = \sum_{l=1}^{r_v + r_w} \bigotimes_{\mu=1}^{d} x_\mu^{(l)},$$

  where

$$x_\mu^{(l)} = \begin{cases} v_\mu^{(l)}, & 1 \le l \le r_v, \\ w_\mu^{(l - r_v)}, & r_v + 1 \le l \le r_v + r_w. \end{cases}$$

- **Matrix-vector multiplication:**

$$\mathbf{A} = \sum_{l_A=1}^{r_A} \bigotimes_{\mu=1}^{d} A_\mu^{(l_A)}, \quad \mathbf{v} = \sum_{l_v=1}^{r_v} \bigotimes_{\mu=1}^{d} v_\mu^{(l_v)} \quad \Rightarrow \quad \mathbf{A}\mathbf{v} = \sum_{l_A=1}^{r_A} \sum_{l_v=1}^{r_v} \bigotimes_{\mu=1}^{d} A_\mu^{(l_A)} v_\mu^{(l_v)}.$$

Since these operations increase the rank, truncation is typically required to approximate the resulting tensor $\mathbf{v}$ with rank r_v by a tensor $\tilde{\mathbf{v}}$ with a lower rank $r_{\tilde{v}}$ by either fixing the rank $r_{\tilde{v}}$ or by fixing ε such that $\|\mathbf{v} - \tilde{\mathbf{v}}\| < \varepsilon$.

3.3 TT format

Another frequently used tensor format is the TT format – TT now stands for *tensor train* instead of the former *tree tensor* –, which can be obtained by successive singular value decompositions. This is a special case of a more general hierarchical format and has been introduced in quantum physics under the name *Matrix Product States* (MPS), see [18] for details. A tensor $\mathbf{v} \in \mathbb{R}^{k_1 \times \cdots \times k_d}$ is decomposed into d component tensors $\mathbf{v}_i$ of at most order three (the first and last are of order two and are often, for the sake of simplicity, considered as tensors of order three with "boundary condition" $\rho_0 = \rho_d = 1$). That is, the entries of $\mathbf{v}$ are given by

$$\mathbf{v}[i_1, \ldots, i_d] = \sum_{k_1=1}^{\rho_1} \cdots \sum_{k_{d-1}=1}^{\rho_{d-1}} \mathbf{v}_1[1, i_1, k_1]\, \mathbf{v}_2[k_1, i_2, k_2]\, \ldots\, \mathbf{v}_{d-1}[k_{d-2}, i_{d-1}, k_{d-1}]\, \mathbf{v}_d[k_{d-1}, i_d, 1].$$

For fixed indices, the component tensors of rank three can be regarded as matrices, which leads to a more compact representation

$$\mathbf{v}[i_1, \ldots, i_d] = V_1[i_1]\, V_2[i_2] \cdots V_{d-1}[i_{d-1}]\, V_d[i_d]$$

and justifies the original name MPS. Here, the numbers ρ_i are the ranks of the TT tensor, resulting in a rank vector $\rho = [\rho_0, \rho_1, \ldots, \rho_d]$ which determines the complexity of the representation. The main advantages of the TT format are its stability from an algorithmic point of view and reasonable computational costs, provided that the ranks of the tensors are small [8]. The basic operations such as addition and matrix-vector multiplication are more complex than in the canonical format and can be found, for example, in [21]. Converting a tensor from the canonical format to the TT format is trivial, but the TT representation requires more memory. Numerical toolboxes for the TT decomposition of tensors and several algorithms for solving linear systems of equations are available online, see e.g. [22].

3.4 Comparison

Complexity-wise, the canonical format would be the ideal candidate for representing tensors since the number of required parameters depends only linearly on the dimension d, the rank r, and the sizes of the individual vector spaces. It turned out, however, that solving even simple problems using the canonical format is hard in practice due to redundancies and instabilities which can lead to numerical problems [8]. The main advantage of the TT format is its structural simplicity, higher-order tensors are reduced to d products of tensors of at most order three. Similar approaches have been known in quantum physics for a long time, the rigorous mathematical analysis, however, is still work in progress (see [8] and references therein). For our purposes, we will rely on the TT format and the TT toolbox developed by Oseledets et al. [22] and implement simple power iteration schemes to solve the resulting eigenvalue problems as we will show in Section 5. Other tensor formats, however, might be advantageous as well for the analysis of the Perron–Frobenius and Koopman operator. This should be investigated further in the future. One drawback of the TT format is that the decomposition depends on the ordering of the dimensions and thus results in different tensor ranks for different orderings.

4 Tensor-based approximation

In this section, we will present a tensor-based reformulation of Ulam's method and EDMD and show that these methods are equivalent to the corresponding vector-based counterparts. The new formulation enables the use of the low-rank tensor approximation approaches described in the previous section. That is, variables can be approximated with different degrees of accuracy.

4.1 Reformulation of Ulam's method

Given a dynamical system $S : \mathcal{X} \to \mathcal{X}$, $\mathcal{X} \subset \mathbb{R}^d$, define a box $\mathcal{B}$ that contains $\mathcal{X}$, i.e.

$$\mathcal{B} = \mathcal{I}_1 \times \cdots \times \mathcal{I}_d \supset \mathcal{X},$$

where each $\mathcal{I}_\mu = [a_\mu, b_\mu] \subset \mathbb{R}$ is an interval. Furthermore, let each $\mathcal{I}_\mu$ be partitioned into k_μ subintervals $\mathcal{I}_\mu^{i_\mu}$, $i_\mu = 1, \ldots, k_\mu$, such that $\mathcal{I}_\mu^{i_\mu} \cap \mathcal{I}_\mu^{j_\mu} = \varnothing$ for $i_\mu \neq j_\mu$. This results in a partitioning of $\mathcal{B}$ into $\hat{k} = \prod_{\mu=1}^{d} k_\mu$ boxes.

Using again standard multi-index notation, we will denote $\mathbf{i} = (i_1, \dots, i_d)$. Equipped with the mapping

$$\mathbf{i} = (i_1, \dots, i_d) \mapsto \hat{\imath} = 1 + \sum_{\mu=1}^{d} \left(\prod_{\nu=1}^{\mu-1} k_\nu \right) (i_\mu - 1), \tag{6}$$

each multi-index $\mathbf{i}$ corresponds to a number $\hat{\imath} \in \{1, \dots, \hat{k}\}$. This induces a canonical numbering of the boxes

$$\mathcal{B}_{\mathbf{i}} = \mathcal{I}_1^{i_1} \times \cdots \times \mathcal{I}_d^{i_d} = \mathcal{B}_{\hat{\imath}}$$

and the entries of tensors $\mathbf{x} \in \mathbb{R}^{k_1 \times \cdots \times k_d}$ such that $\mathbf{x}_{\mathbf{i}} = x_{\hat{\imath}}$, where $x = \mathrm{vec}(\mathbf{x})$. Thus, with the aid of Ulam's method we could now generate the finite-dimensional representation of the Perron–Frobenius operator $\hat{P} \in \mathbb{R}^{\hat{k} \times \hat{k}}$ as described in Section 2. Our goal, however, is to approximate the operator by a tensor $\mathbf{P} \in \mathbb{R}^{k_1 \times \cdots \times k_d \times k_1 \times \cdots \times k_d}$. Note that the indicator function for the box $\mathcal{B}_{\mathbf{i}}$ can be written as

$$\mathbb{1}_{\mathcal{B}_{\mathbf{i}}}(x) = \prod_{\mu=1}^{d} \mathbb{1}_{\mathcal{I}_\mu^{i_\mu}}(x_\mu) = \mathbb{1}_{\mathcal{B}_{\hat{\imath}}}(x). \tag{7}$$

That is, each d-dimensional indicator function $\mathbb{1}_{\mathcal{B}_{\mathbf{i}}}(x)$ is now written as a product of d one-dimensional indicator functions $\mathbb{1}_{\mathcal{I}_\mu^{i_\mu}}(x_\mu)$.

Example 4.1. Let us start with a simple example which illustrates the idea behind the tensor-based formulation. Consider the box discretization $\{\mathcal{B}_1, \dots, \mathcal{B}_9\}$ of $\mathcal{B} = [0, 3]^2$ shown in Figure 2(a). Thus, using Ulam's method, we would obtain 9 indicator functions $\{\mathbb{1}_{\mathcal{B}_1}, \dots, \mathbb{1}_{\mathcal{B}_9}\}$ and the matrix $\hat{P}$ that approximates the Perron–Frobenius operator $\mathcal{P}$ would be a row-stochastic (9×9)-matrix. An example of such a matrix is shown in Figure 2(b), the underlying dynamical system is not relevant here. The goal now is to rewrite this matrix using tensors (cf. Example 4.3).

(a)

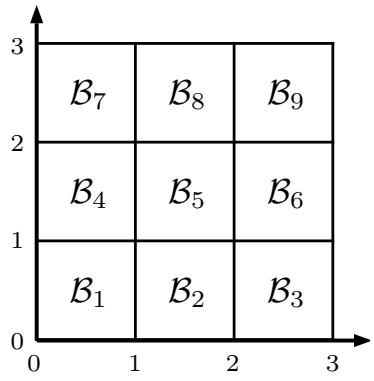

(b)

$$\hat{P} = \begin{bmatrix} 0.68 & 0.09 & 0.07 & 0.04 & 0.02 & 0 & 0.09 & 0.01 & 0 \\ 0.36 & 0.06 & 0.40 & 0.03 & 0 & 0.02 & 0.05 & 0.03 & 0.05 \\ 0.07 & 0.12 & 0.64 & 0.02 & 0 & 0.07 & 0.03 & 0 & 0.05 \\ 0.31 & 0.07 & 0.02 & 0.09 & 0.01 & 0.02 & 0.39 & 0.05 & 0.04 \\ 0.25 & 0.05 & 0.18 & 0.06 & 0.01 & 0.02 & 0.17 & 0.05 & 0.21 \\ 0.06 & 0.06 & 0.37 & 0.01 & 0 & 0.04 & 0.07 & 0.03 & 0.36 \\ 0.17 & 0.01 & 0.01 & 0.09 & 0.01 & 0.01 & 0.60 & 0.05 & 0.05 \\ 0.05 & 0 & 0.06 & 0.08 & 0 & 0.05 & 0.29 & 0.13 & 0.34 \\ 0.01 & 0 & 0.03 & 0.02 & 0.02 & 0.11 & 0.05 & 0.09 & 0.67 \end{bmatrix}$$

Figure 2: (a) Box discretization of $\mathcal{B} = [0, 3]^2$. (b) Example of a resulting approximation $\hat{P}$ of the Perron–Frobenius operator $\mathcal{P}$, obtained by applying Ulam's method.

Defining intervals $\mathcal{I}_\mu^1 = [0, 1]$, $\mathcal{I}_\mu^2 = [1, 2]$, and $\mathcal{I}_\mu^3 = [2, 3]$ as well as indicator functions

$$\mathbb{1}_{\mathcal{I}_\mu^{i_\mu}}(x_\mu) = \begin{cases} 1, & x_\mu \in \mathcal{I}_\mu^{i_\mu}, \\ 0, & \text{otherwise}, \end{cases}$$

for $\mu = 1, 2$ and $i_\mu = 1, 2, 3$, the ansatz functions for the box discretization can be written as

$$\begin{aligned} \mathbb{1}_{\mathcal{B}_{1,1}}(x) &= \mathbb{1}_{\mathcal{I}_1^1}(x_1)\mathbb{1}_{\mathcal{I}_2^1}(x_2) = \mathbb{1}_{\mathcal{B}_1}(x), \\ \mathbb{1}_{\mathcal{B}_{2,1}}(x) &= \mathbb{1}_{\mathcal{I}_1^2}(x_1)\mathbb{1}_{\mathcal{I}_2^1}(x_2) = \mathbb{1}_{\mathcal{B}_2}(x), \\ &\vdots \qquad\qquad \vdots \\ \mathbb{1}_{\mathcal{B}_{2,3}}(x) &= \mathbb{1}_{\mathcal{I}_1^2}(x_1)\mathbb{1}_{\mathcal{I}_2^3}(x_2) = \mathbb{1}_{\mathcal{B}_8}(x), \\ \mathbb{1}_{\mathcal{B}_{3,3}}(x) &= \mathbb{1}_{\mathcal{I}_1^3}(x_1)\mathbb{1}_{\mathcal{I}_2^3}(x_2) = \mathbb{1}_{\mathcal{B}_9}(x). \end{aligned}$$

△

The product formulation of the indicator functions naturally leads to a tensor approximation $\mathbf{P}$ of the Perron–Frobenius operator $\mathcal{P}$. Let $Q_\mu : \mathbb{R}^d \to \mathbb{R}$ be the projection onto the μ-th component of a vector, i.e. $Q_\mu(x) = x_\mu$. Then we define

$$\mathbf{P}[i_1, \dots, i_d, j_1, \dots, j_d] = \frac{1}{n} \sum_{l=1}^{n} \prod_{\mu=1}^{d} \mathbb{1}_{\mathcal{I}_\mu^{j_\mu}} \left(Q_\mu \left(S \left(x_{i_1,\dots,i_d}^{(l)} \right) \right) \right), \tag{8}$$

where $x_{i_1,\dots,i_d}^{(l)}$, $l = 1, \dots, n$, are the test points generated for box $\mathcal{B}_{\mathbf{i}}$. Instead of checking to which d-dimensional box the test points are mapped, the d dimensions are now treated separately.

Proposition 4.2. *It holds that*

$$\mathbf{Pv} = \lambda \mathbf{v} \quad \Leftrightarrow \quad \hat{P} v = \lambda v.$$

Proof. It suffices to show that $\mathbf{P}[i_1, \dots, i_d, j_1, \dots, j_d] = \hat{P}_{\hat{i}\hat{j}}$ and that $(\mathbf{Pv})[i_1, \dots, i_d] = (\hat{P}v)_{\hat{i}}$. The entries of $\mathbf{P}$ and $\hat{P}$ are identical since with (2) and (7)

$$\begin{aligned}
\hat{P}_{\hat{i}\hat{j}} &= \frac{1}{n} \sum_{l=1}^{n} \mathbb{1}_{\mathcal{B}_{\hat{j}}} \left(S \left(x_{\hat{i}}^{(l)} \right) \right) \\
&= \frac{1}{n} \sum_{l=1}^{n} \prod_{\mu=1}^{d} \mathbb{1}_{\mathcal{I}_\mu^{j_\mu}} \left(Q_\mu \left(S \left(x_{i_d,\dots,i_d}^{(l)} \right) \right) \right) = \mathbf{P}[i_1, \dots, i_d, j_1, \dots, j_d].
\end{aligned}$$

For $d = 1$, the multi-index $\mathbf{i}$ is mapped to $\hat{i} = i_1$ and $\mathbf{j}$ to $\hat{j} = j_1$ by (6). Furthermore, $\hat{k} = k_1$ and

$$(\mathbf{Pv})[i_1] = \sum_{j_1=1}^{k_1} \mathbf{P}[i_1, j_1] \mathbf{v}[j_1] = \sum_{\hat{j}=1}^{\hat{k}} \hat{P}_{\hat{i}\hat{j}} v_{\hat{j}} = (\hat{P}v)_{\hat{i}}.$$

Then, we obtain by induction

$$\begin{aligned}
(\mathbf{Pv})[i_1, \dots, i_{d+1}] &= \sum_{j_1=1}^{k_1} \cdots \sum_{j_{d+1}=1}^{k_{d+1}} \mathbf{P}[i_1, \dots, i_{d+1}, j_1, \dots, j_{d+1}] \mathbf{v}[j_1, \dots, j_{d+1}] \\
&= \sum_{j_{d+1}=1}^{k_{d+1}} \left(\sum_{j_1=1}^{k_1} \cdots \sum_{j_d=1}^{k_d} \mathbf{P}[i_1, \dots, i_d, i_{d+1}, j_1, \dots, j_d, j_{d+1}] \mathbf{v}[j_1, \dots, j_d, j_{d+1}] \right) \\
&= \sum_{j_{d+1}=1}^{k_{d+1}} \left(\sum_{j_1=1}^{k_1} \cdots \sum_{j_d=1}^{k_d} \mathbf{P}^{(i_{d+1}, j_{d+1})}[i_1, \dots, i_d, j_1, \dots, j_d] \mathbf{v}^{(j_{d+1})}[j_1, \dots, j_d] \right) \\
&= \begin{bmatrix} \hat{P}^{(i_{d+1},1)} & \dots & \hat{P}^{(i_{d+1},k_{d+1})} \end{bmatrix} \begin{bmatrix} v^{(1)} \\ \vdots \\ v^{(k_{d+1})} \end{bmatrix} \begin{matrix} \\ \\ \} \in \mathbb{R}^{k_1 \cdots k_d} \end{matrix} \\
&= (\hat{P}v)_{\hat{i}}.
\end{aligned}$$

Here, the matrices and vectors with the superscripts (i_{d+1}, j_{d+1}) and (j_{d+1}), respectively, are obtained by fixing the corresponding indices, that is, these matrices and vectors are lower-dimensional slices of the corresponding higher-dimensional objects. Since the entries of $v^{(j_{d+1})}$ are indexed by $(j_1, \dots, j_d) \mapsto \hat{j} = 1 + \sum_{\mu=1}^{d} \left(\prod_{\nu=1}^{\mu-1} k_\nu \right) (j_\mu - 1)$, the entries of the larger vector v – obtained by stacking the vectors $v^{(j_{d+1})} \in \mathbb{R}^{k_1 \cdots k_d}$ – are indexed by

$$(j_1, \dots, j_{d+1}) \mapsto \hat{j} = 1 + \sum_{\mu=1}^{d} \left(\prod_{\nu=1}^{\mu-1} k_\nu \right) (j_\mu - 1) + (k_1 \cdots k_d)(j_{d+1} - 1) = 1 + \sum_{\mu=1}^{d+1} \left(\prod_{\nu=1}^{\mu-1} k_\nu \right) (j_\mu - 1). \qquad \square$$

Analogously, left eigenvectors of $\mathbf{P}$ correspond to left eigenvectors of $\hat{P}$. The implementation of the tensor-based formulation is straightforward since only index computations for intervals are required, a numbering of the d-dimensional boxes is not needed anymore. Let $\mathbb{T}$ be the set of all test points and $\mathrm{ind} : \mathbb{R}^d \to \mathbb{N}^d$ the functions that returns the corresponding multi-index $\mathbf{i}$ for a point $x \in \mathbb{R}^d$ so that $x_\mu \in I_\mu^{i_\mu}$, $\mu = 1, \ldots, d$. Then, Ulam's method can simply be expressed as:

for each test point $x \in \mathbb{T}$ **do**
  $y = S(x)$
  $\mathbf{i} = \mathrm{ind}(x)$
  $\mathbf{j} = \mathrm{ind}(y)$
  $\mathbf{P}[i_1, \ldots, i_d, j_1, \ldots, j_d] \leftarrow \mathbf{P}[i_1, \ldots, i_d, j_1, \ldots, j_d] + \frac{1}{n}$
end for

In the standard formulation, the rows of the matrix $\hat{P}$ sum up to one. Correspondingly, the sum of all entries of each subtensor of $\mathbf{P}$ with fixed multi-index $\mathbf{i}$ is one, i.e.

$$\sum_{j_1=1}^{k_1} \cdots \sum_{j_d=1}^{k_d} \mathbf{P}[i_1, \ldots, i_d, j_1, \ldots, j_d] = 1.$$

Example 4.3. Let us consider Example 4.1 again. We select n random test points $x_{i_1,i_2}^{(l)}$ for each box $\mathcal{B}_{i_1,i_2}$, $i_1, i_2 = 1, \ldots, 3$ and $l = 1, \ldots, n$. This leads to a new approximation $\mathbf{P} \in \mathbb{R}^{3\times3\times3\times3}$. Written in Matlab notation, we would obtain

$$\mathbf{P}[:,:,1,1] = \begin{bmatrix} 0.68 & 0.31 & 0.17 \\ 0.36 & 0.25 & 0.05 \\ 0.07 & 0.06 & 0.01 \end{bmatrix}, \quad \mathbf{P}[:,:,1,2] = \begin{bmatrix} 0.04 & 0.09 & 0.09 \\ 0.03 & 0.06 & 0.08 \\ 0.02 & 0.01 & 0.02 \end{bmatrix}, \quad \mathbf{P}[:,:,1,3] = \begin{bmatrix} 0.09 & 0.39 & 0.60 \\ 0.05 & 0.17 & 0.29 \\ 0.03 & 0.07 & 0.05 \end{bmatrix},$$

$$\mathbf{P}[:,:,2,1] = \begin{bmatrix} 0.09 & 0.07 & 0.01 \\ 0.06 & 0.05 & 0 \\ 0.12 & 0.06 & 0 \end{bmatrix}, \quad \mathbf{P}[:,:,2,2] = \begin{bmatrix} 0.02 & 0.01 & 0.01 \\ 0 & 0.01 & 0 \\ 0 & 0 & 0.02 \end{bmatrix}, \quad \mathbf{P}[:,:,2,3] = \begin{bmatrix} 0.01 & 0.05 & 0.05 \\ 0.03 & 0.05 & 0.13 \\ 0 & 0.03 & 0.09 \end{bmatrix},$$

$$\mathbf{P}[:,:,3,1] = \begin{bmatrix} 0.07 & 0.02 & 0.01 \\ 0.40 & 0.18 & 0.06 \\ 0.64 & 0.37 & 0.03 \end{bmatrix}, \quad \mathbf{P}[:,:,3,2] = \begin{bmatrix} 0 & 0.02 & 0.01 \\ 0.02 & 0.02 & 0.05 \\ 0.07 & 0.04 & 0.11 \end{bmatrix}, \quad \mathbf{P}[:,:,3,3] = \begin{bmatrix} 0 & 0.04 & 0.05 \\ 0.05 & 0.21 & 0.34 \\ 0.05 & 0.36 & 0.67 \end{bmatrix}.$$

Note that each matrix $\mathbf{P}[:,:,j_1,j_2]$ corresponds to a column of matrix $\hat{P}$ in Figure 2. For the resulting eigenvalue problem, we obtain – using a simple power iteration, see Section 5 – the left eigenvector $\mathbf{v}_1$ corresponding to the largest eigenvalue $\lambda_1 = 1$

$$\mathbf{v}_1 = \begin{bmatrix} 0.6503 & 0.1393 & 0.4501 \\ 0.1046 & 0.0261 & 0.0901 \\ 0.4355 & 0.0864 & 0.3719 \end{bmatrix},$$

which is a good approximation of the largest left eigenvector v_1 of the matrix $\hat{P}$ given by

$$v_1 = \begin{bmatrix} 0.6503 & 0.1393 & 0.4501 & 0.1046 & 0.0261 & 0.0901 & 0.4355 & 0.0864 & 0.3719 \end{bmatrix}. \qquad \triangle$$

Instead of working with the full format, the matrix $\mathbf{P}$ can also be expressed directly using the canonical tensor format. Assume that a test point x is mapped from box $\mathcal{B}_\mathbf{i}$ to box $\mathcal{B}_\mathbf{j}$, where $\mathbf{i} = (i_1, \ldots, i_d)$ and $\mathbf{j} = (j_1, \ldots, j_d)$ are again multi-indices. Now let $e_\mu^{i_\mu} \in \mathbb{R}^{k_\mu}$ be the i_μ-th unit vector of size k_μ and let

$$\mathbf{e}^\mathbf{i} = \bigotimes_{\mu=1}^{d} e_\mu^{i_\mu}.$$

Then the elementary tensor $\mathbf{e}^\mathbf{i} \otimes \mathbf{e}^\mathbf{j} \in \mathbb{R}^{k_1\times\cdots\times k_d\times k_1\times\cdots\times k_d}$ describes the mapping of this point. Furthermore, let ind be again the function that returns the multi-index of the box that contains the point x and $\mathbb{T}$ the set of all test points, then the matrix $\mathbf{P}$ can be represented as a sum of elementary tensors of this form, i.e.

$$\mathbf{P} = \frac{1}{n} \sum_{x\in\mathbb{T}} \mathbf{e}^{\mathrm{ind}(x)} \otimes \mathbf{e}^{\mathrm{ind}(S(x))}.$$

That is, the number of elementary tensors is $\hat{k}n$, i.e. the number of boxes multiplied by the number of test points per box, and thus potentially too large to store (unless sparse tensor formats are used). However, we will store only a low-rank approximation to reduce the required storage space. Note that this elementary tensor representation can also be easily converted into the TT format for numerical computations using the TT toolbox.

The question now is whether the tensor representation offers advantages over the standard formulation of Ulam's method. The goal is to approximate the eigenfunctions of the Perron–Frobenius operator or Koopman operator using low-rank tensors, reducing the computational cost as well as the memory consumption. Before we present numerical results, let us also rewrite EDMD in tensor form.

4.2 Reformulation of EDMD

Instead of writing Ψ as a vector of functions $\Psi = [\psi_1, \psi_2, \dots, \psi_{\hat{k}}]^T$, we now write Ψ as a tensor of functions. We start by selecting basis functions for each dimension separately. Let

$$\mathcal{D}_\mu = \{\psi_\mu^1, \dots, \psi_\mu^{k_\mu}\}$$

be the set of basis functions for dimension μ, $\mu = 1, \dots, d$. Here, each $\psi_\mu^{i_\mu} : \mathbb{R} \to \mathbb{R}$ depends only on x_μ. Then our tensor basis for EDMD contains all functions of the form

$$\boldsymbol{\psi}^{\mathbf{i}}(x) = \prod_{\mu=1}^{d} \psi_\mu^{i_\mu}(x_\mu), \tag{9}$$

where $\mathbf{i} = (i_1, \dots, i_d)$ is a multi-index. Thus,

$$\mathcal{D} = \left\{ \prod_{\mu=1}^{d} \psi_\mu^{i_\mu}, \psi_\mu^{i_\mu} \in \mathcal{D}_\mu \right\}.$$

That is, we have again $\hat{k} = \prod_{\mu=1}^{d} k_\mu$ basis functions and $\boldsymbol{\Psi} : \mathbb{R}^d \to \mathbb{R}^{k_1 \times \cdots \times k_d}$, with

$$\boldsymbol{\Psi}[i_1, \dots, i_d](x) = \boldsymbol{\psi}^{\mathbf{i}}(x).$$

Example 4.4. Let us begin with a simple example: Assume we have a two-dimensional domain $\mathcal{X} \subset \mathbb{R}^2$ and we want to use monomials of order up to three $\{1, x_\mu, x_\mu^2, x_\mu^3\}$ in x_1 and x_2 direction to approximate the eigenfunctions of the Koopman operator. Written in tensor form, we obtain

$$\boldsymbol{\Psi}(x) = \begin{bmatrix} 1 & x_2 & x_2^2 & x_2^3 \\ x_1 & x_1x_2 & x_1x_2^2 & x_1x_2^3 \\ x_1^2 & x_1^2x_2 & x_1^2x_2^2 & x_1^2x_2^3 \\ x_1^3 & x_1^3x_2 & x_1^3x_2^2 & x_1^3x_2^3 \end{bmatrix}.$$

That is, $\boldsymbol{\Psi}[i_1, i_2](x) = x_1^{i_1-1}x_2^{i_2-1}$. Analogously, for a d-dimensional domain, we would obtain $\boldsymbol{\Psi}(x) \in \mathbb{R}^{k_1 \times \cdots \times k_d}$ with

$$\boldsymbol{\Psi}[i_1, \dots, i_d](x) = x_1^{i_1-1} \cdots x_d^{i_d-1}. \qquad \triangle$$

Such a tensor basis is often used for high-dimensional problems, see also [23]. Typical basis functions are monomials, Hermite polynomials, or trigonometric functions. In the standard formulation, all basis functions are enumerated and rewritten in vector form (3). The difference here is that the tensor form will be preserved. We will use again (6) as a mapping from multi-index to single index when required.

Now $\mathbf{A}, \mathbf{G} \in \mathbb{R}^{k_1 \times \cdots \times k_d \times k_1 \times \cdots \times k_d}$ can be constructed as follows:

$$\begin{aligned} \mathbf{A}[i_1, \dots, i_d, j_1, \dots, j_d] &= \langle \mathcal{K}\boldsymbol{\Psi}[i_1, \dots, i_d],\, \boldsymbol{\Psi}[j_1, \dots, j_d] \rangle, \\ \mathbf{G}[i_1, \dots, i_d, j_1, \dots, j_d] &= \langle \boldsymbol{\Psi}[i_1, \dots, i_d],\, \boldsymbol{\Psi}[j_1, \dots, j_d] \rangle. \end{aligned}$$

The entries are again – as in the standard EDMD formulation – approximated using a collocation approach. EDMD computes the entries as shown in (4), for the new tensor-based formulation this results in

$$\mathbf{A}[i_1,\dots,i_d,j_1,\dots,j_d] = \frac{1}{m}\sum_{l=1}^{m} \mathbf{\Psi}[i_1,\dots,i_d](y_l)\mathbf{\Psi}[j_1,\dots,j_d](x_l),$$
$$\mathbf{G}[i_1,\dots,i_d,j_1,\dots,j_d] = \frac{1}{m}\sum_{l=1}^{m} \mathbf{\Psi}[i_1,\dots,i_d](x_l)\mathbf{\Psi}[j_1,\dots,j_d](x_l),$$

or in short form, using the outer product,

$$\begin{aligned}\mathbf{A} &= \frac{1}{m}\sum_{l=1}^{m} \mathbf{\Psi}(y_l)\otimes\mathbf{\Psi}(x_l),\\ \mathbf{G} &= \frac{1}{m}\sum_{l=1}^{m} \mathbf{\Psi}(x_l)\otimes\mathbf{\Psi}(x_l),\end{aligned} \tag{10}$$

which in turn results in a generalized eigenvalue problem of the form

$$\boldsymbol{\xi}\mathbf{A} = \lambda\boldsymbol{\xi}\mathbf{G}.$$

Note that the eigenvalue problem is the same as (5) in the standard case. For the sake of simplicity, we are omitting the index $\mathbf{i}$ here.

Proposition 4.5. *Provided that the basis functions can be written in tensor product form (9),*

$$\boldsymbol{\xi}\mathbf{A} = \lambda\boldsymbol{\xi}\mathbf{G} \quad\Leftrightarrow\quad \xi\hat{A} = \lambda\xi\hat{G}.$$

Proof. We just show that the entries of $\mathbf{A}$ and $\hat{A}$ as well as the entries of $\mathbf{G}$ and $\hat{G}$ are identical, the rest – the equivalency of the matrix-vector and tensor products – follows from Proposition 4.2. Assuming that the basis can be written in the product form, we obtain from (4)

$$\begin{aligned}\hat{a}_{\hat{i}\hat{j}} &= \frac{1}{m}\sum_{l=1}^{m}\psi_{\hat{i}}(y_l)\psi_{\hat{j}}(x_l)\\ &= \frac{1}{m}\sum_{l=1}^{m}\prod_{\mu=1}^{d}\psi_\mu^{i_\mu}(Q_\mu(y_l))\prod_{\mu=1}^{d}\psi_\mu^{j_\mu}(Q_\mu(x_l))\\ &= \frac{1}{m}\sum_{l=1}^{m}\mathbf{\Psi}[i_1,\dots,i_d](y_l)\mathbf{\Psi}[j_1,\dots,j_d](x_l)\\ &= \mathbf{A}[i_1,\dots,i_d,j_1,\dots,j_d]\end{aligned}$$

and analogously $\hat{g}_{\hat{i}\hat{j}} = \mathbf{G}[i_1,\dots,i_d,j_1,\dots,j_d]$. Here, Q_μ is again the projection onto the μ-th component of a vector, cf. (8). □

Instead of storing the dense matrices $\mathbf{A}$ and $\mathbf{G}$, we can again directly represent these matrices using the canonical tensor format. The basis was chosen in such a way that $\mathbf{\Psi}(x)$ can be written as

$$\mathbf{\Psi}(x) = \bigotimes_{\mu=1}^{d}\tilde{\psi}_\mu(x_\mu),$$

where $\tilde{\psi}_\mu = [\psi_\mu^1,\dots,\psi_\mu^{k_\mu}]^T \in \mathbb{R}^{k_\mu}$. With (10) it follows that $\mathbf{A}$ and $\mathbf{G}$ can be written as sums of m elementary tensors. As before, we are not storing the full-rank tensor, but only low-rank approximations.

The eigentensors $\boldsymbol{\xi} \in \mathbb{R}^{k_1\times\cdots\times k_d}$ of the generalized eigenvalue problem can then be used to approximate the eigenfunctions of the Perron–Frobenius operator or Koopman operator: Let $\boldsymbol{\xi}$ be a left eigentensor, then

$$\varphi(x) = \langle\boldsymbol{\xi},\, \mathbf{\Psi}(x)\rangle$$

approximates an eigenfunction of the Koopman operator. Analogously, if $\boldsymbol{\xi}$ is a right eigentensor of the generalized eigenvalue problem – observe that $\mathbf{G}[i_1, \ldots, i_d, j_1, \ldots, j_d] = \mathbf{G}[j_1, \ldots, j_d, i_1, \ldots, i_d]$ –, then $\varphi(x)$ is an approximation of the corresponding eigenfunction of the Perron–Frobenius operator, see also [12].

To compute the dominant eigenfunctions of the Koopman operator or Perron–Frobenius operator, we will use simple power iteration schemes outlined in the next section. General purpose eigenvalue solvers for nonsymmetric generalized eigenvalue problems are, to our knowledge, not part of the tensor libraries yet. Solvers for symmetric (non-generalized) eigenvalue problems already exist and are part of the TT toolbox [22].

5 Eigenvalue problems

In Section 2 and Section 4, we have shown that in order to compute the eigenfunctions of the Perron–Frobenius operator and Koopman operator, respectively, using either Ulam's method or EDMD, we need to solve standard eigenvalue problems or generalized eigenvalue problems. For the reformulated version of these methods, we have to develop the required numerical algorithms to solve the resulting tensor-based eigenvalue problems. At the time of writing, we are not aware of any tensor toolbox containing numerical methods for nonsymmetric generalized eigenvalue problems. Methods for the computation of eigenvectors of symmetric positive definite matrices in the TT format have been proposed in [8], where the eigenvalue problem is rewritten as a (Rayleigh quotient based) minimization problem which is then solved using the *Alternating Linear Scheme* (ALS). In practice, these methods have recently also been successfully used for nonsymmetric problems, although convergence has not been shown yet [24].

Suitable methods for eigenvalue problems can be subdivided into two main categories as explained in [10] (see also references therein, e.g. [9]): The first category of methods is based on combining classical iterative algorithms with low-rank truncation after each step, the second is based on a reformulation as an optimization problem, where admissible solutions are constrained to the set of low-rank tensors. In this section, we will describe a generalization of simple power iteration methods – belonging to the first category – to tensor-based eigenvalue problems. For a detailed description of general power iteration methods, we refer to [25]. Power iteration and inverse power iteration for tensors have also been proposed in [9]. The main difference between the standard algorithms and the tensor-based counterpart is that for the latter truncation is used to keep the ranks of the tensors low. It is important that the iteration moves from the initial state to the final state without creating intermediate solutions with an excessive rank [9].

5.1 Power iteration methods for standard eigenvalue problems

In what follows, let $\mathcal{T}$ denote the truncation of a tensor. Then instead of a classical iteration scheme of the form $x_{k+1} = F(x_k)$, we simply obtain $\mathbf{x}_{k+1} = \mathcal{T}(F(\mathbf{x}_k))$. For an eigenvalue problem of the form $\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$, given an initial guess $\mathbf{v}_0$ for the dominant eigenvector, the power iteration algorithm computes:

for $k = 1, 2, \ldots$ **do**
$\mathbf{w}^{(k)} = \mathcal{T}\left(\mathbf{A}\mathbf{v}^{(k-1)}\right)$
$\mathbf{v}^{(k)} = \mathbf{w}^{(k)} / \left\|\mathbf{w}^{(k)}\right\|$
$\lambda^{(k)} = \left\langle \mathbf{v}^{(k)}, \mathcal{T}(\mathbf{A}\mathbf{v}^{(k)}) \right\rangle$
end for

The iteration converges to an eigenvector associated with the largest eigenvalue λ_1 of the truncated operator $\mathbf{A}$ if the eigenvalue is simple and the initial guess $\mathbf{v}^{(0)}$ has a component in the direction of the corresponding dominant eigenvector $\mathbf{v}_1$ [25]. Even if the initial guess does not have a component in the direction of $\mathbf{v}_1$, rounding errors typically ensure that this direction will be picked up during the iteration. The rate of convergence depends on the ratio between the second-largest and largest eigenvalue λ_2/λ_1. The main advantage of this method is that it requires only matrix-vector multiplications and can thus easily be used for tensor eigenvalue problems.

A modification of this algorithm to compute eigenvectors corresponding to any eigenvalue is the inverse power iteration with shift, which – assuming that $\mathbf{A} - \theta\mathbf{I}$ is nonsingular – can be written in the form:

for $k = 1, 2, \ldots$ **do**
Solve $(\mathbf{A} - \theta\mathbf{I})\mathbf{w}^{(k)} = \mathbf{v}^{(k-1)}$
$\mathbf{v}^{(k)} = \mathbf{w}^{(k)} / \left\|\mathbf{w}^{(k)}\right\|$
$\lambda^{(k)} = \left\langle \mathbf{v}^{(k)}, \mathbf{A}\mathbf{v}^{(k)} \right\rangle$

end for

The parameter θ is called shift and the iteration converges to the eigenvalue closest to θ. This method is just the standard power iteration applied to the matrix $(\mathbf{A} - \theta\mathbf{I})^{-1}$. Here, the linear solver computes a low-rank approximation of the solution so that truncation is not required.

5.2 Power iteration methods for generalized eigenvalue problems

Given matrices $\mathbf{A}$ and $\mathbf{B}$, the generalized eigenvalue problem is given by $\mathbf{Av} = \lambda\mathbf{Bv}$. In this case, the power iteration method also requires the solution of a linear system of equations. Thus, we can also directly apply the inverse power iteration, where the resulting systems of linear equations are again solved with ALS:

for $k = 1, 2, \ldots$ **do**
  Solve $(\mathbf{A} - \theta\mathbf{B})\mathbf{w}^{(k)} = \mathbf{Bv}^{(k-1)}$
  $\mathbf{v}^{(k)} = \mathbf{w}^{(k)} / \left\|\mathbf{w}^{(k)}\right\|$
  $\lambda^{(k)} = \left\langle \mathbf{v}^{(k)}, \mathbf{Av}^{(k)} \right\rangle / \left\langle \mathbf{v}^{(k)}, \mathbf{Bv}^{(k)} \right\rangle$
end for

In order to keep the ranks of the intermediate solutions low, we also approximate the matrices $\mathbf{P}$ (Ulam's method) or $\mathbf{A}$ and $\mathbf{G}$ (EDMD) by low-rank tensors. That is, the initial matrices are converted to $\tilde{\mathbf{P}}$, $\tilde{\mathbf{A}}$, and $\tilde{\mathbf{G}}$ with a lower rank since the rank of these matrices can initially be very high. We are typically only interested in the general behavior of the eigenfunctions, a highly accurate representation of the eigenfunctions is often not needed.

6 Examples

All examples presented within this section have been implemented in Matlab using – for the sake of efficiency – `mex`-functions to integrate the SDEs. All tensor computations were carried out with the *TT toolbox* [22]. For the eigenvector computations, we used our implementation of the simple power iteration methods described in Section 5.

6.1 2-dimensional double well problem

Let us start with a simple 2-dimensional example, a stochastic differential equation of the form

$$\begin{aligned} dx_1 &= -\nabla_{x_1} V(x_1, x_2)\, dt + \sigma\, dW_1, \\ dx_2 &= -\nabla_{x_2} V(x_1, x_2)\, dt + \sigma\, dW_2, \end{aligned}$$

where W_1 and W_2 are two independent standard Wiener processes. Here, the potential is given by

$$V(x_1, x_2) = (x_1^2 - 1)^2 + x_2^2,$$

see Figure 3 (cf. [12]). Furthermore, we set $\sigma = 0.7$. Note that in this case the potential can be written as $V(x_1, x_2) = V_1(x_1) + V_2(x_2)$. In order to analyze the tensor-based methods, we rotate the potential by an angle α and obtain

$$\widetilde{V}(x_1, x_2) = \left((\cos(\alpha)\, x_1 - \sin(\alpha)\, x_2)^2 - 1\right)^2 + (\sin(\alpha)\, x_1 + \cos(\alpha)\, x_2)^2 .$$

The two independent Wiener processes W_1 and W_2 are rotated accordingly. We would expect that the eigenfunctions of systems with small α can be accurately approximated by low-rank tensors, whereas systems with a larger value of α require higher ranks since the dynamics are not aligned with the axes anymore.

The second eigenfunctions of the Perron–Frobenius operator for the systems with potential $\widetilde{V}$ and different values of α computed using the tensor-based version of Ulam's method are shown in Figure 4. We chose $\alpha = 0$, $\alpha = \pi/12$, $\alpha = \pi/6$, and $\alpha = \pi/4$. The domain $\mathcal{X} = [-2, 2]^2$ was subdivided into 50×50 equally sized boxes. That is, $\mathbf{P} \in \mathbb{R}^{50\times50\times50\times50}$. For each box, 100 randomly chosen test points were generated. The Euler–Maruyama method with a step size $h = 10^{-3}$ was used for the numerical integration, where one evaluation of S corresponds to $10,000$ integration steps, that is, the integration interval is $[0, 10]$. The shift parameter θ of the power iteration method was set to a value slightly smaller than 1. Figure 5 illustrates how the tensor approximation, depending on the rank, successively picks up the information about the shape of the eigenfunction and generates more and more accurate representations. For $\alpha = \pi/4$, a tensor of rank 1 cannot represent the minimum and maximum

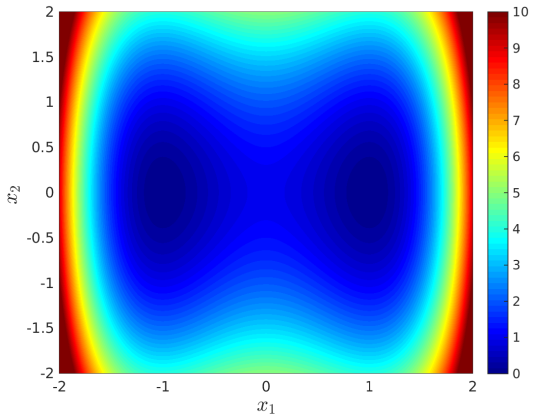

Figure 3: Double-well potential $V(x_1, x_2) = (x_1^2 - 1)^2 + x_2^2$.

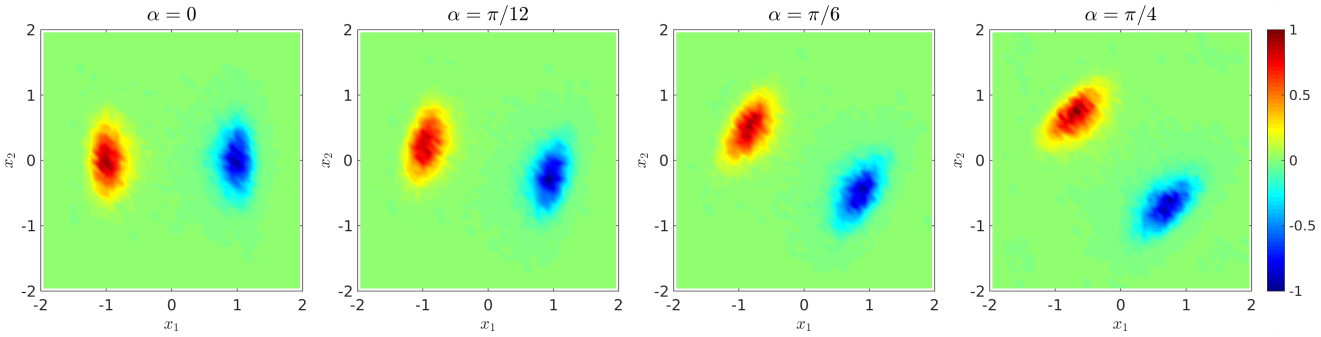

Figure 4: Second eigenfunction of the Perron–Frobenius operator for different values of α.

simultaneously since this would lead to two additional peaks in the lower left and upper right corner. Here, the first pair of singular vectors represents the maximum, the second pair of singular vectors the minimum.

Additionally, we computed the eigenfunctions with the standard version of Ulam's method to evaluate the accuracy of the approximation and compared it with the results obtained by using the new tensor-based formulation. Figure 6 shows the influence of the truncation of the operator as well as the influence of the truncation of the resulting eigenfunctions. Here, in order to analyze the accuracy, we also compare the first eigenfunction with the analytically computed invariant density. Since we are computing eigenfunctions of the Perron–Frobenius operator associated with a stochastic differential equation, the results depend strongly on the number of test points chosen for each box. The higher the number of test points per box, the smoother the eigenfunction approximation. Thus, in this case, the smoother low-rank solutions can counterintuitively lead to better approximations of the true eigenfunctions. The high ranks are mainly required to resolve the numerical noise introduced by the coarse approximation of the operator. This can be seen, for example, in Figure 6b. Decreasing the rank initially reduces the error – the truncation of the operator results in smoother eigenfunctions – until the shape of the eigenfunction cannot be described by a low-rank approximation anymore and the error increases. For $\alpha = 0$, the x_1 and x_2 dynamics are independent and a low-rank approximation is sufficient. Furthermore, the results illustrate that for a fixed-rank approximation, the error is smaller when the system's dynamics are aligned with the axes.

This example shows that in order to be able to approximate eigenfunctions by low-rank tensors, the dynamics of the system should be aligned with the axes chosen, although even if the dynamics are not aligned, the tensor format might be advantageous. In general, the dynamics are unknown a priori and not necessarily aligned with the axes, but for higher-dimensional systems it is often possible to decompose a system into slow and fast subsystems. Not all variables of a system might be equally important to describe the system's behavior. The intuition would be that certain subsystems require less information and that the tensor approximation automatically captures the relevant dynamics, using high ranks only when necessary.

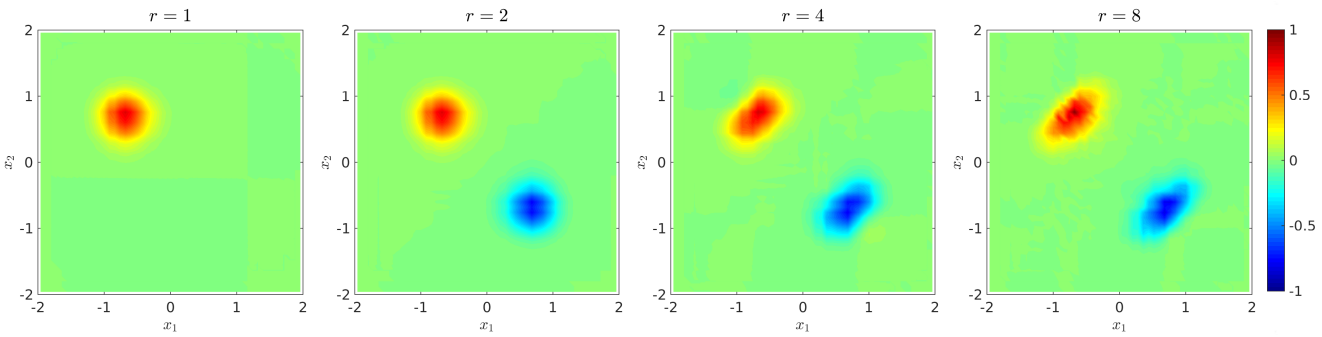

Figure 5: Tensor approximations of the second eigenfunction of the Perron–Frobenius operator for $\alpha = \pi/4$ with increasing rank r.

6.2 3-dimensional triple well problem

Let us consider a more complex 3-dimensional example with the potential function

$$V(x_1, x_2, x_3) = 3e^{-x_1^2-(x_2-\frac{1}{3})^2} - 3e^{-x_1^2-(x_2-\frac{5}{3})^2} - 5e^{-(x_1-1)^2-x_2^2} - 5e^{-(x_1+1)^2-x_2^2} + \tfrac{2}{10}x_1^4 + \tfrac{2}{10}\left(x_2 - \tfrac{1}{3}\right)^4 + x_3^2.$$

This is a potential taken from [20], augmented by a third dimension. We subdivided the domain $\mathcal{X} = [-2, 2] \times [-1, 2] \times [-2, 2]$ into $20 \times 20 \times 20$ boxes of the same size. For each box, we randomly generated 1000 test points. The resulting finite-dimensional approximation is then a tensor $\mathbf{P} \in \mathbb{R}^{20\times20\times20\times20\times20\times20}$. Figure 7 shows a scatter plot of the first three eigenfunctions, which were computed using the inverse power iteration described in Section 5. The first eigenfunction clearly shows the three expected regions with high probabilities corresponding to the minima of the potential V. The second eigenfunction separates the two deeper wells at $(-1, 0, 0)$ and $(1, 0, 0)$, the third eigenfunction separates these wells from the third shallower well at $(0, 1.5, 0)$. The differences between the eigenfunctions computed using the conventional and the tensor-based version of Ulam's method are negligible, the average difference between the first eigenvector v_1 and the tensor $\mathbf{v}_1$ is of the order of 10^{-6}, which is mainly due to the less accurate power iteration method applied to the tensor eigenvalue problem.

For the sake of comparison, we computed the eigenfunctions of the Koopman operator using EDMD. We chose basis functions $\mathcal{D} = \{x_1^{i_1} x_2^{i_2} x_3^{i_3},\ i_1, i_2, i_3 = 0, \ldots, 5\}$. Hence, $\mathbf{A}, \mathbf{G} \in \mathbb{R}^{6\times6\times6\times6\times6\times6}$. For higher-order monomials, the resulting matrices are ill-conditioned and the eigenfunctions cannot be computed accurately anymore. The results are shown in Figure 8. As before, the second eigenfunction separates the two deeper wells. The third eigenfunction separates these wells from the third shallower well and is close to zero for the regions around the deep wells. The trivial eigenfunction of the Koopman operator corresponding to $\lambda_1 = 1$ is not plotted here since it is constant and does not contain relevant information about the system. Note that the eigenvalues are slightly different due to the different set of basis functions used for EDMD.

7 Conclusion

We have reformulated the problems of computing finite-dimensional approximations of the Perron–Frobenius and Koopman operator in a different format using tensors instead of vectors. The matrices $\mathbf{P}$ (if Ulam's method is used) or $\mathbf{A}$ and $\mathbf{G}$ (if EDMD is used) can now either be assembled in the dense tensor format or directly in the canonical tensor format – which can then easily be converted into the TT format –, enabling low-rank approximations of the aforementioned operators. The next step is to systematically develop the numerical methods required to efficiently solve the resulting nonsymmetric generalized tensor eigenvalue problems and also to store and handle these tensors minimizing memory requirements so that even high-dimensional problems can be solved. First results obtained by applying simple algorithms such as power iteration methods are promising and show that the approaches presented within this paper might be able to tackle high-dimensional problems. Currently, several toolboxes for tensor-based problems are under development. Once these toolboxes contain methods for solving nonsymmetric generalized eigenvalue problems, the proposed approaches can be implemented easily, potentially facilitating the computation of meta-stable sets or almost invariant sets of dynamical systems that could not be handled before due to the

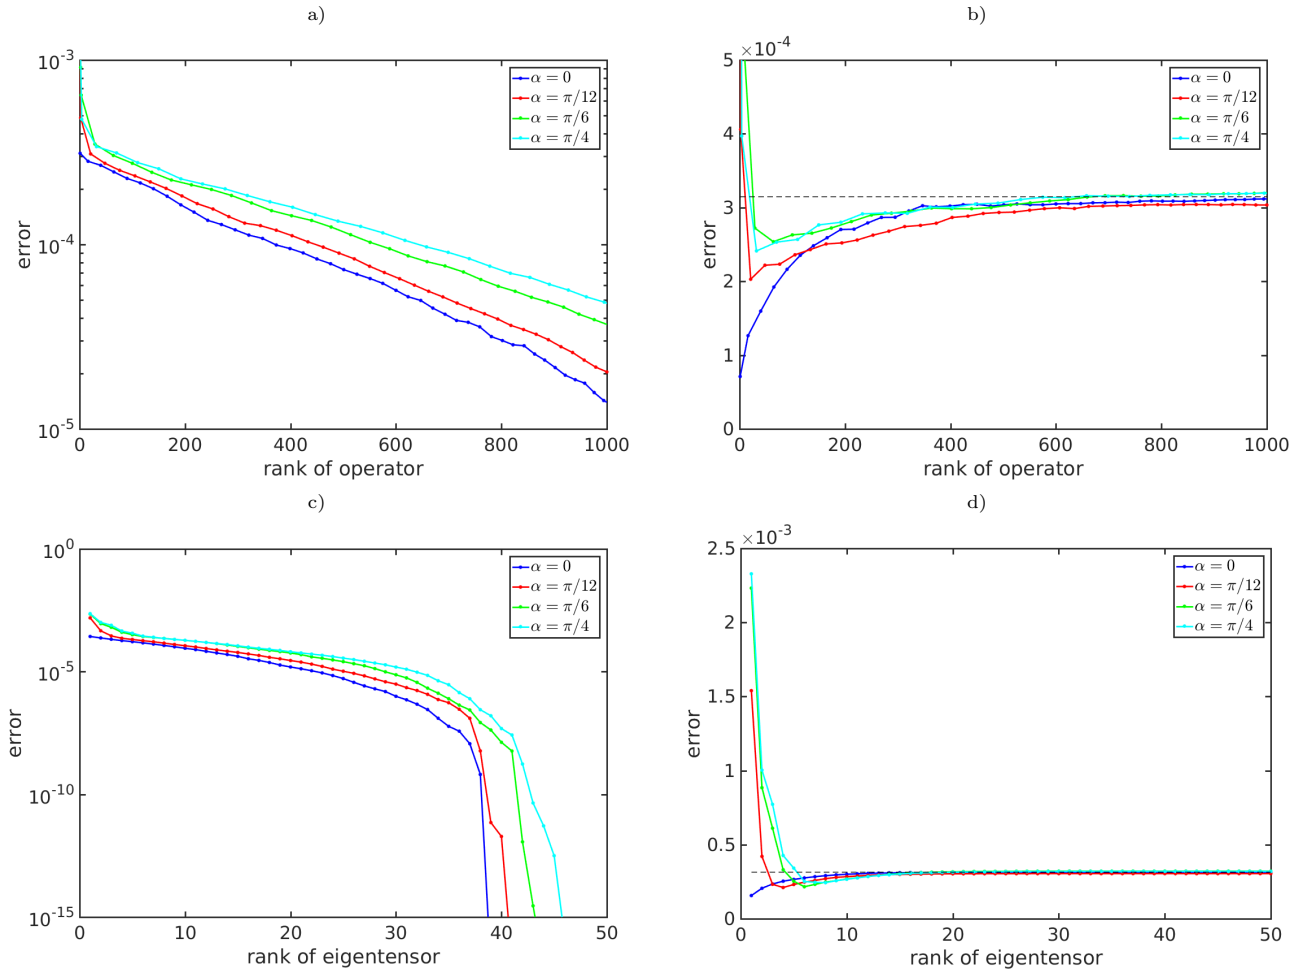

Figure 6: (Top) Let $\hat{v}_1$ denote the first eigenvector of $\hat{P}$, v_1 the (vectorized) eigentensor of the truncated tensor representation $\mathbf{P}$, and μ_{inv} the analytically computed invariant density. The error here is defined by $e = \frac{1}{k} \|v_1 - \hat{v}_1\|_2$ and $e = \frac{1}{k} \|v_1 - \mu_{\text{inv}}\|_2$, respectively, where k is the number of boxes. a) Difference between v_1 and $\hat{v}_1$ depending on the rank of $\mathbf{P}$. b) Difference between v_1 and μ_{inv}. (Bottom) Influence of the truncation of the first eigentensor v_1 of the full tensor representation $\mathbf{P}$ on the accuracy. c) Difference between the truncated eigentensor v_1 and $\hat{v}_1$. d) Difference between the truncated eigentensor v_1 and μ_{inv}. The dashed lines show the error for the full-rank approximation which is almost identical for the different values of α.

curse of dimensionality. We demonstrated the tensor-based version of Ulam's method and EDMD using two- and three-dimensional problems mainly because the results can be easily validated and visualized.

Future work also includes determining which tensor format is suited best for our purposes. Currently, one of the main bottlenecks is the simulation required to obtain the data. Even simple molecular dynamics simulations, for instance, might easily take several days, but in order to capture the behavior of the system, long trajectories or a large number of short simulations with different initial conditions are required. This huge amount of data must then be processed. Thus, also the construction of the matrices $\mathbf{P}$ or $\mathbf{A}$ and $\mathbf{G}$ is time-consuming. The number of boxes or basis functions required to represent each variable of the system accurately is in general unknown a priori. If we are, for instance, only interested in the leading eigenfunctions of the Koopman operator, the fast variables of the system are typically almost constant and require less information to be captured. Starting with a set of only a few basis functions for each unknown, which is then, if needed, augmented adaptively based on the system's behavior would greatly improve the efficiency. Adaptive methods combined with (sparse) tensor approaches might be able to tackle high-dimensional systems and diminish the curse of dimensionality. Furthermore, a detailed numerical analysis of the efficiency and accuracy of the proposed algorithms would help understand the limitations and find

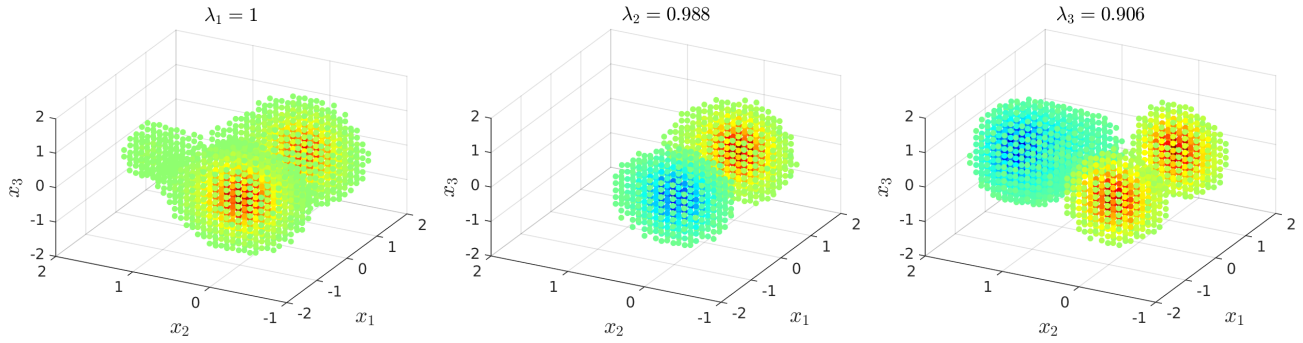

Figure 7: Scatter plot of the first three eigenfunctions of the Perron–Frobenius operator for the triple well system. Only entries whose absolute value is larger than a given threshold are plotted, entries close to zero are omitted.

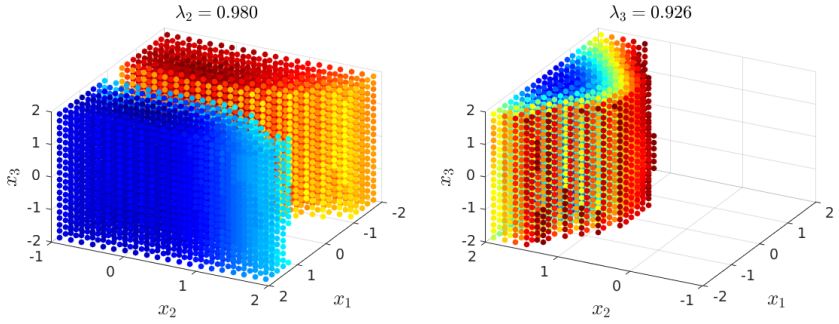

Figure 8: Scatter plot of the second and third eigenfunction of the Koopman operator for the triple well system. Note that compared to the other plots the second eigenfunction of the Koopman operator is rotated by 180 degrees around the x_3 axis for a better visualization.

opportunities for improvement.

Another open problem is the – depending on the number of dimensions d – typically extremely large condition number of the matrices $\mathbf{A}$ or $\mathbf{G}$ if EDMD with, for instance, monomials are used to compute the eigenfunctions of the Perron–Frobenius or Koopman operator. Hence, the resulting eigenvalue problems cannot be solved accurately anymore for high-dimensional systems. A detailed understanding and numerical analysis of different basis functions might help mitigate this problem. Radial-basis functions or other more locally defined functions could lead to better results.

Acknowledgements

This research has been partially funded by Deutsche Forschungsgemeinschaft (DFG) through grant CRC 1114. Moreover, we would like to thank the reviewers for their helpful comments and suggestions.

References

[1] M. Dellnitz and O. Junge. On the approximation of complicated dynamical behavior. *SIAM Journal on Numerical Analysis*, 36(2):491–515, 1999.

[2] M. Dellnitz, G. Froyland, and O. Junge. The algorithms behind GAIO – Set oriented numerical methods for dynamical systems. In *Ergodic theory, analysis, and efficient simulation of dynamical systems*, pages 145–174. Springer, 2000.

[3] R. Preis, M. Dellnitz, M. Hessel, Ch. Schütte, and E. Meerbach. Dominant paths between almost invariant sets of dynamical systems. DFG Schwerpunktprogramm 1095, Technical Report 154, 2004.

[4] M. Budišić, R. Mohr, and I. Mezić. Applied Koopmanism. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(4), 2012.

[5] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley. A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *ArXiv e-prints*, 2014.

[6] G. Froyland, G. Gottwald, and A. Hammerlindl. A computational method to extract macroscopic variables and their dynamics in multiscale systems. *SIAM Journal on Applied Dynamical Systems*, 13(4):1816–1846, 2014.

[7] G. Froyland, G. A. Gottwald, and A. Hammerlindl. A trajectory-free framework for analysing multiscale systems. *ArXiv e-prints*, 2014.

[8] S. Holtz, T. Rohwedder, and R. Schneider. The alternating linear scheme for tensor optimization in the tensor train format. *SIAM Journal on Scientific Computing*, 34(2):A683–A713, 2012.

[9] G. Beylkin and M. J. Mohlenkamp. Algorithms for numerical analysis in high dimensions. *SIAM Journal on Scientific Computing*, 26(6):2133–2159, 2005.

[10] L. Grasedyck, D. Kressner, and C. Tobler. A literature survey of low-rank tensor approximation techniques. *ArXiv e-prints*, 2013.

[11] Jiu Ding, Q. Du, and Tien-Yien Li. High order approximation of the Frobenius–Perron operator. *Applied Mathematics and Computation*, 53:151–171, 1993.

[12] S. Klus, P. Koltai, and Ch. Schütte. On the numerical approximation of the Perron–Frobenius and Koopman operator. *ArXiv e-prints*, 2015.

[13] M. O. Williams, C. W. Rowley, and I. G. Kevrekidis. A kernel-based approach to data-driven Koopman spectral analysis. *ArXiv e-prints*, 2014.

[14] G. Chen and T. Ueta, editors. *Chaos in Circuits and Systems*. World Scientific Series on Nonlinear Science, Series B, Volume 11. World Scientific, 2002.

[15] E. M. Bollt and N. Santitissadeekorn. *Applied and Computational Measurable Dynamics*. Society for Industrial and Applied Mathematics, 2013.

[16] F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. J. S. Mey, and F. Noé. Variational approach to molecular kinetics. *Journal of Chemical Theory and Computation*, 10(4):1739–1752, 2014.

[17] F. Nüske, R. Schneider, F. Vitalini, and F. Noé. Variational tensor approach for approximating the rare-event kinetics of macromolecular systems. *The Journal of Chemical Physics*, 144(5), 2016.

[18] W. Hackbusch. Numerical tensor calculus. *Acta Numerica*, 23:651–742, 2014.

[19] G. Beylkin and M. J. Mohlenkamp. Numerical operator calculus in higher dimensions. *Proceedings of the National Academy of Sciences*, 99(16):10246–10251, 2002.

[20] Ch. Schütte and M. Sarich. *Metastability and Markov State Models in Molecular Dynamics: Modeling, Analysis, Algorithmic Approaches*. Number 24 in Courant Lecture Notes. American Mathematical Society, 2013.

[21] I. V. Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.

[22] I. V. Oseledets. *TT-toolbox 2.0: Fast multidimensional array operations in TT format*, 2011.

[23] G. Friesecke, O. Junge, and P. Koltai. Mean field approximation in conformation dynamics. *Multiscale Modeling & Simulation*, 8(1):254–268, 2009.

[24] P. Gelß, S. Matera, and Ch. Schütte. Solving the master equation without kinetic Monte Carlo: Tensor train approximations for a CO oxidation model. *Journal of Computational Physics*, 314:489–502, 2016.

[25] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 4th edition, 2013.